

THE DISTRIBUTION OF TORSION SUBSCHEMES OF ABELIAN VARIETIES

JEFFREY D. ACHTER

ABSTRACT. We consider the distribution of p -power group schemes among the torsion of abelian varieties over finite fields of characteristic p , as follows. Fix natural numbers g and n , and let ξ be a non-supersingular principally quasipolarized Barsotti-Tate group of level n . We classify the \mathbb{F}_q -rational forms ξ^α of ξ . Among all principally polarized abelian varieties X/\mathbb{F}_q of dimension g with $X[p^n]_{\mathbb{F}_q} \cong \xi_{\mathbb{F}_q}$, we compute the frequency with which $X[p^n] \cong \xi^\alpha$. The error in our estimate is bounded by D/\sqrt{q} , where D depends on g, n and p , but not on ξ .

1. INTRODUCTION

Either the general philosophy of arithmetic statistics or the specific question of Cohen-Lenstra heuristics for function fields can easily lead one to wonder:

How are group schemes distributed among the N -torsion group schemes of principally polarized abelian varieties over a finite field?

More concretely, for various combinations of $g, q = p^e$, and N , we would like to understand the multiset

$$\{(X[N], \lambda_N) : (X, \lambda) \in \mathcal{A}_g(\mathbb{F}_q)\}.$$

For example, suppose N is relatively prime to p and g is fixed. The N -torsion of the universal abelian scheme $\mathcal{X}[N] \rightarrow \mathcal{A}_g$ is a local system of (\mathbb{Z}/N) -modules, and the scheme of isomorphisms $\text{Isom}((\mathcal{X}[N], \lambda_N), ((\mathbb{Z}/N)^{\oplus 2g}, \langle \cdot, \cdot \rangle))$ is an étale, Galois cover of \mathcal{A}_g . A monodromy calculation shows that this cover is irreducible; the Chebotarev theorem then gives good control, as $q \rightarrow \infty$, of the distribution of forms of $(\mathbb{Z}/N)^{\oplus 2g}$. (See, e.g., [1].)

The landscape looks dramatically different if, instead of considering prime-to- p torsion, one analyzes p^n -torsion. For example, there are several different geometric isomorphism classes of group scheme which arise as $X[p]$. Perhaps worse yet, there are infinitely many different geometric isomorphism classes of group schemes which arise as $X[p^2]$ over $\overline{\mathbb{F}}_p$. Therefore, a direct strategy based on local systems and the Chebotarev theorem seems difficult to implement.

Still, not all is lost. In [5], the authors fix a base field \mathbb{F}_q and describe a natural measure on the space of principally quasipolarized group schemes of rank $2g$ which are killed by p . They then are able to compute the large- g expectation of various events, such as having specified a -number or p -rank. Unfortunately, it is unknown whether this measure accurately describes the distribution of the p -torsion group schemes of abelian varieties, although numerical experiments with hyperelliptic Jacobians are encouraging.

In the present paper, we take a complementary, “fixed g , large q ” approach. An initial, naïve attempt fails to detect many of the interesting structures on \mathcal{A}_g in positive characteristic. The ordinary locus is open and dense, and so in the large q limit, (essentially) all $X[p]$ are forms of $(\mu_p \oplus \mathbb{Z}/p)^g$. To detect more subtle variations, we *condition* on geometric isomorphism class, as follows.

This work was partially supported by grants from the Simons Foundation (204164) and the NSA (H98230-14-1-0161).

Let ζ be a geometric isomorphism class of principally quasipolarized BT_n ; let $\mathcal{A}_{g,\zeta}$ be the locus where $\mathcal{X}[p^n]$ is geometrically isomorphic to ζ . We will show that there is a finite group $A(\zeta)$ whose conjugacy classes $A(\zeta)^\sharp$ parametrize, for suitable fields \mathbb{F}_q , \mathbb{F}_q -rational forms of ζ . We essentially show (see Theorem 4.5 for details) that, for each $\alpha \in A(\zeta)^\sharp$, the proportion of elements of $(X, \lambda) \in \mathcal{A}_{g,\zeta}(\mathbb{F}_q)$ for which $(X, \lambda)[p^n] \cong \zeta^\alpha$ approaches $1/\#Z_{A(\zeta)}(\alpha)$ in a surprisingly uniform way. In fact,

$$(1.1) \quad \frac{\#\{(X, \lambda) \in \mathcal{A}_{g,\zeta}(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \zeta^\alpha\}}{\#\mathcal{A}_{g,\zeta}(\mathbb{F}_q)} = \frac{1}{\#Z_{A(\zeta)}(\alpha)} + O_{g,n,p}(1/\sqrt{q});$$

the error is bounded by D/\sqrt{q} for some constant D which depends only on g , n and p . Thanks to Lemma 2.2, this is compatible with the broad philosophy of [5], and even that of [9]; a (polarized) group (scheme) occurs with frequency inversely proportional to the size of its automorphism group.

Our analysis starts (Section 2) with a classification of the possible \mathbb{F}_q -rational forms of a given principally quasipolarized BT_n ζ . One fact which emerges is that there is a finite set which controls such forms, independent of the base field. In this way, it is possible to compare forms of ζ over ever-larger finite extensions of the base field.

In Section 3 we review (and, where necessary, extend) results, largely due to Oort and to Chai and Oort, on the structure of the strata $\mathcal{A}_{g,\zeta}$. It turns out that, although there are infinitely many such strata (if $n \geq 2$), their topology (if $n \gg_g 0$) may be uniformly bounded in terms of g , n and p .

We apply these results in Section 4 to state and prove our main equidistribution result. As is often the case, a geometric Chebotarev theorem is the crucial guarantor of equidistribution, but it is invoked here in a somewhat novel way. If n is small relative to g , then there is often no single local system on $\mathcal{A}_{g,\zeta}$ which controls the variation in the forms of ζ . Still, we express $\mathcal{A}_{g,\zeta}$ as a disjoint union of infinitely many central leaves $\mathcal{A}_{g,\tilde{\zeta}}$. (Of course, only finitely many of them are defined over a given finite field.) Since each of *these* has bounded topology, we are able to deduce the uniform bound asserted in (1.1).

The paper concludes (Section 5) with a series of examples, and a brief discussion of related results.

The present investigation was inspired by the work of Cais, Ellenberg and Zureick-Brown, even though the results here are incomparable with those of [5]. The author acknowledges helpful discussions with Oort and Liedtke, especially concerning the example in Section 5.3. In particular, the explicit classification of Dieudonné modules there was communicated to the author by Oort. Finally, this paper benefited from the referee's careful reading.

2. FORMS OF A BT_n

2.1. Notation. We will work over finite fields of characteristic $p > 0$; in particular, q is always a power of p . Let \mathbb{F} be an algebraic closure of \mathbb{F}_p .

Let K be a perfect field of characteristic p , and let n be a natural number. A Barsotti-Tate group of level n , or BT_n , over K is a commutative finite group scheme G/K which is annihilated by p^n and flat over \mathbb{Z}/p^n . (If $n = 1$, one further insists that $\mathrm{im} F : G \rightarrow G^{(p)} = \ker V : G^{(p)} \rightarrow G$.) A Barsotti-Tate, or p -divisible, group is an inductive limit $G = \varinjlim G_n$ where G_n is a BT_n and $[p]$ is an epimorphism; we will sometimes refer to such an object as a BT_∞ .

Let G be a BT_n . The Cartier dual of G is $G^D = \mathrm{Hom}(G, G_m)$. If $p > 2$, a principal quasipolarization on G is an isomorphism $\lambda : G \xrightarrow{\sim} G^D$ such that $\lambda^D = -\lambda \circ \kappa$, where $\kappa : G^{DD} \xrightarrow{\sim} G$ is the canonical isomorphism. (In characteristic 2, this definition must be modified slightly; see [22, 2.6] or [14, 6.3])

for more details.) The pair (G, λ) is called a principally quasipolarized, or pqp, group scheme (or BT_n).

Define:

- $\Xi_{g,n}(\mathbb{F})$ for $1 \leq n \leq \infty$ is the set of \mathbb{F} -isomorphism classes of principally quasipolarized BT_n 's of height $2g$ over \mathbb{F} .
- For $\xi \in \Xi_{g,n}(\mathbb{F})$, let $\xi(\mathbb{F}_q)$ be the set of \mathbb{F}_q -isomorphism classes of pqp group schemes $(G, \lambda)/\mathbb{F}_q$ such that the class of $(G, \lambda)_{\mathbb{F}_q}$ in $\Xi(\mathbb{F})$ is ξ .

For each of these sets we will write, e.g., $(G, \lambda) \in \xi(\mathbb{F}_q)$ (or even $G \in \xi(\mathbb{F}_q)$) when we really mean a representative for the \mathbb{F}_q -isomorphism class. For $1 \leq n \leq m \leq \infty$ there is an obvious truncation functor $\Xi_{g,m}(\mathbb{F}) \rightarrow \Xi_{g,n}(\mathbb{F})$, which we will denote by $\tilde{\xi} \mapsto \tilde{\xi}[p^n]$. Let $\Xi_{g,m}(\mathbb{F})_{\tilde{\xi}}$ be the fiber of this map over $\tilde{\xi}$.

For an abstract group Γ , let $\Gamma^{\#}$ be the set of conjugacy classes of Γ , and let $Z(\Gamma)$ be the center of Γ .

2.2. Automorphism groups and forms. Let K be a finite field of characteristic p , let n be a positive integer, and let $(G, \lambda)/K$ be a pqp- BT_n . Its automorphism group scheme $\text{Aut}_{G,\lambda}$ is affine, and sits in the usual étale-connected sequence

$$(2.1) \quad 0 \longrightarrow \text{Aut}_{G,\lambda}^0 \longrightarrow \text{Aut}_{G,\lambda} \longrightarrow \pi_0(\text{Aut}_{G,\lambda}) \longrightarrow 0.$$

Then $\pi_0(\text{Aut}_{G,\lambda})(\bar{K})$ is a $\text{Gal}(K)$ -module. We will find it convenient to call (G, λ) split if the group $\pi_0(\text{Aut}_{G,\lambda})(\bar{K})$ has trivial $\text{Gal}(K)$ -structure, i.e., if $\pi_0(\text{Aut}_{G,\lambda})(K) = \pi_0(\text{Aut}_{G,\lambda})(\bar{K})$. For a split pqp group scheme $(G, \lambda)/K$, let $A_{G,\lambda}$ be the abstract, finite group $\pi_0(\text{Aut}_{G,\lambda})(K)$.

Since K is perfect, (2.1) splits canonically, and we sometimes view $A_{G,\lambda}$ as a (finite, discrete) group of automorphisms of (G, λ) .

For a finite extension L/K , let $\text{Twist}(L/K, (G, \lambda))$ be the set of K -isomorphism classes of pqp group schemes $(H, \mu)/K$ such that $(H, \mu)_L \cong (G, \lambda)_L$. On one hand, $\text{Twist}(L/K, (G, \lambda))$ is a pointed set, with distinguished element the K -isomorphism class of (G, λ) itself. On the other hand, we endow $A_{G,\lambda}^{\#}$ with the structure of a pointed set by distinguishing (the conjugacy class of) the identity element. Let $\text{Twist}(K, (G, \lambda)) = \cup_{L/K} \text{Twist}(L/K, (G, \lambda))$.

Lemma 2.1. *Let $(G, \lambda)/K$ be a split pqp group scheme over a finite field. There is a bijection of pointed sets between*

- $\text{Twist}(K, (G, \lambda));$ and
- $A_{G,\lambda}^{\#}$.

Proof. It suffices to prove that for each finite extension L/K , there is a bijection of pointed sets between $\text{Twist}(L/K, (G, \lambda))$ and conjugacy classes of elements of $A_{G,\lambda}$ of order dividing $[L : K]$. By [26, III.1.3], there is a bijection between $\text{Twist}(L/K, (G, \lambda))$ and $H^1(L/K, \text{Aut}_{G,\lambda}(L))$. Consider the exact sequence of groups

$$0 \longrightarrow \text{Aut}_{G,\lambda}^0(L) \longrightarrow \text{Aut}_{G,\lambda}(L) \longrightarrow \pi_0(\text{Aut}_{G,\lambda})(L) \longrightarrow 0$$

and the corresponding long exact sequence of Galois cohomology groups

$$H^1(L/K, \text{Aut}_{G,\lambda}^0(L)) \longrightarrow H^1(L/K, \text{Aut}_{G,\lambda}(L)) \longrightarrow H^1(L/K, \pi_0(\text{Aut}_{G,\lambda})(L)) \longrightarrow H^2(L/K, \text{Aut}_{G,\lambda}^0(L)).$$

The left-most term is zero by Lang's theorem; the rightmost term is trivial, since K is a finite field and thus has cohomological dimension ≤ 1 . Since G is split, $\pi_0(\text{Aut}_{G,\lambda})(L) \cong A_{G,\lambda}$, and thus $H^1(L/K, \text{Aut}_{G,\lambda}(L)) \cong H^1(L/K, A_{G,\lambda})$. Moreover, $\text{Gal}(L/K)$ is cyclic and $A_{G,\lambda}$ is trivial as $\text{Gal}(L/K)$ -module, whence the asserted description of $H^1(L/K, A_{G,\lambda})$. \square

In particular, (the isomorphism class of) a twist $(H, \mu) \in \text{Twist}(K, (G, \lambda))$ is represented by some $\alpha \in A_{G, \lambda}$, well-defined up to conjugacy.

Lemma 2.2. *Suppose $H \in \text{Twist}(K, (G, \lambda))$ is represented by $\alpha \in A_{G, \lambda}$.*

(a) *We have*

$$\pi_0(\text{Aut}_{H, \mu})(K) \cong Z_{A_{G, \lambda}}(\alpha).$$

(b) *If L is a finite extension of K , then $H_L \in \text{Twist}(L, (G, \lambda)_L)$ is represented by $\alpha^{[L:K]} \in A_{(G, \lambda)_L} \cong A_{G, \lambda}$.*

Proof. Part (a) is a special case of a more general statement worked out in [26, I.5.3.34]; part (b) follows immediately from the fact that the $[L : K]^{\text{th}}$ power of a topological generator of $\text{Gal}(K)$ is a topological generator of $\text{Gal}(L)$. \square

In particular, in the context of Lemma 2.2, (H, μ) is split if and only if $\alpha \in Z(A_{G, \lambda})$. We will call such an (H, μ) a central twist of (G, λ) .

Lemma 2.3. *Let $\xi \in \Xi_{g, n}(\mathbb{F})$. Suppose $(G, \lambda) \in \xi(\mathbb{F}_{q^r})$ and $(H, \mu) \in \xi(\mathbb{F}_{q^s})$ are split and that $\gcd(r, s) = 1$. Then there are split $(G', \lambda') \in \xi(\mathbb{F}_{q^r})$ and $(H', \mu') \in \xi(\mathbb{F}_{q^s})$ such that $(G', \lambda')_{\mathbb{F}_{q^{rs}}} \cong (H', \mu')_{\mathbb{F}_{q^{rs}}}$.*

Proof. Since $(G, \lambda)_{\mathbb{F}_{q^{rs}}}$ and $(H, \mu)_{\mathbb{F}_{q^{rs}}}$ are both split, $(H, \mu)_{\mathbb{F}_{q^{rs}}}$ is a central twist of $(G, \lambda)_{\mathbb{F}_{q^{rs}}}$; there is some $a \in Z(A_{G, \lambda}) \subset A_{G, \lambda}$ such that twisting $(G, \lambda)_{\mathbb{F}_{q^{rs}}}$ by the cocycle $\sigma \mapsto a$ yields a pqp group scheme isomorphic to $(H, \mu)_{\mathbb{F}_{q^{rs}}}$. A split twist (G', μ') of $(G, \mu)_{\mathbb{F}_{q^s}}$ corresponds to some $b \in Z(A_{G, \lambda})$. Using (in this lemma only) additive notation for $Z(A_{G, \lambda})$, we find that $(G', \mu')_{\mathbb{F}_{q^{rs}}}$ corresponds to sb (Lemma 2.2(b)). Similarly, if (H', μ') is a split twist of (H, μ) corresponding to $c \in Z(A_{H, \mu})$, then $(H', \mu')_{\mathbb{F}_{q^{rs}}}$ is represented by sc as a twist of $(H, \mu)_{\mathbb{F}_{q^{rs}}}$, and by $rc + a$ as a twist of $(G, \lambda)_{\mathbb{F}_{q^{rs}}}$ [26, I.5.3.35]. Consequently, in order to prove the lemma, we must show that for any $a \in Z(A_{G, \lambda})$, there exist $b, c \in Z(A_{G, \lambda})$ such that

$$rc + a = sb.$$

If we write $Z(A_{G, \lambda}) \cong \oplus \mathbb{Z}/m_i$ and recall that r and s are relatively prime, the existence of b and c is obvious. \square

Proposition 2.4. *Let $\xi \in \Xi_{g, n}(\mathbb{F})$. There exists a finite field $\mathbb{F}_{q(\xi)}$ and a split pqp group scheme $(G, \lambda) \in \xi(\mathbb{F}_{q(\xi)})$ such that, if $(H, \mu) \in \xi(\mathbb{F}_q)$ is split, then \mathbb{F}_q contains $\mathbb{F}_{q(\xi)}$ and (H, μ) is a central twist of $(G, \lambda)_{\mathbb{F}_q}$.*

Proof. Given Lemma 2.3, it suffices to prove that if $\gcd(r, s) = 1$, and if there are split pqp group schemes $(G_r, \lambda_r) \in \xi(\mathbb{F}_{q^r})$ and $(G_s, \lambda_s) \in \xi(\mathbb{F}_{q^s})$ which are isomorphic to the pqp group scheme $(G, \lambda) \in \xi(\mathbb{F}_{q^{rs}})$, then (G, λ) descends to \mathbb{F}_q . Let σ generate $\text{Gal}(\mathbb{F}_{q^{rs}}/\mathbb{F}_q)$; by hypothesis, (G, λ) is canonically isomorphic to its pullbacks by σ^s and σ^r , and thus is invariant under $\text{Gal}(\mathbb{F}_{q^{rs}}/\mathbb{F}_q)$. Descent is effective for affine group schemes, and the proposition is proved. \square

Set

$$A(\xi) = A_{G, \lambda},$$

where $(G, \lambda) \in \xi(\mathbb{F}_q)$ is split. As above, we note that $A(\xi)$ acts on (G, λ) .

If $\alpha \in A(\xi)$ and if \mathbb{F}_q contains $\mathbb{F}_{q(\xi)}$, let

$$\xi^{\alpha, \mathbb{F}_q} \in \xi(\mathbb{F}_q)$$

denote the form of ξ corresponding to the conjugacy class of α in $A(\xi)$.

This notation is slightly ambiguous; the isomorphism class of $\zeta^{\alpha, \mathbb{F}_q}$ actually depends on the choice of split model (G, λ) for ζ . Still, suppose that (G_1, λ_1) and (G_2, λ_2) are split models for ζ . Recall (as seen in Lemma 2.2) that, after a choice of isomorphism $A_{G_1, \lambda_1} \rightarrow A_{G_2, \lambda_2}$ of abstract groups, the bijection between $\text{Twist}(\mathbb{F}_q, (G_1, \lambda_1))$ and $\text{Twist}(\mathbb{F}_q, (G_2, \lambda_2))$ is realized by

$$A_{G_1, \lambda_1}^{\#} \longrightarrow A_{G_2, \lambda_2}^{\#}$$

$$[\alpha] \longmapsto [\gamma\alpha]$$

where γ is some element of the center. Now suppose that (H, μ) is an \mathbb{F}_q -form of ζ , represented by $[\alpha_i]$ in $\text{Twist}(\mathbb{F}_q, (G_i, \lambda_i))$. Since the centralizers of α_1 and $\gamma\alpha_1$ coincide, the quantity

$$\#Z_{A_{(G_i, \lambda_i)}}(\alpha_i)$$

is independent of the choice of split model for ζ .

2.3. Lifts. Suppose $1 \leq n \leq m \leq \infty$ and $\tilde{\zeta} \in \Xi_{g, m}(\mathbb{F})$; let $\zeta = \tilde{\zeta}[p^n] \in \Xi_{g, m}(\mathbb{F})$ be its n -truncation. We define a subgroup $A(\tilde{\zeta}, \zeta) \subseteq A(\zeta)$ as follows.

Choose a split representative $(\tilde{G}, \tilde{\lambda})/\mathbb{F}_q$ for $\tilde{\zeta}$, and let $(G, \lambda) = (\tilde{G}[p^n], \tilde{\lambda}[p^n])$. The natural morphism of functors induces a map of group schemes

$$\text{Aut}_{\tilde{G}, \tilde{\lambda}} \xrightarrow{\rho} \text{Aut}_{G, \lambda}$$

and in particular a map of étale quotients

$$\pi_0(\text{Aut}_{\tilde{G}, \tilde{\lambda}}) \xrightarrow{\pi_0(\rho)} \pi_0(\text{Aut}_{G, \lambda}).$$

Let

$$A(\tilde{\zeta}, \zeta) = \text{im}(\pi_0(\rho)).$$

As an abstract group, it is independent of the choice of model $(\tilde{G}, \tilde{\lambda})$, and we view it as the group of discrete automorphisms of ζ which lift to $\tilde{\zeta}$.

3. STRATIFICATIONS ON \mathcal{A}_g

Let $N \geq 3$ be an integer relatively prime to p . Let $\mathcal{A}_{g, d, N}$ be the moduli space, over \mathbb{F}_p , of g -dimensional abelian varieties equipped with a degree d^2 polarization and symplectic level- N structure, and let

$$\mathcal{A}_g = \mathcal{A}_{g, 1, N}.$$

Let $\mathcal{X} \rightarrow \mathcal{A}_g$ be the universal principally polarized abelian scheme (with level N structure).

Lemma 3.1. *Fix $\zeta \in \Xi_{n, g}(\mathbb{F})$ with $1 \leq n \leq \infty$. There exists a reduced subscheme $\mathcal{A}_{g, \zeta} \subset \mathcal{A}_g$ such that, for each algebraically closed field k of characteristic p ,*

$$(3.1) \quad \mathcal{A}_{g, \zeta}(k) = \{(X, \lambda) \in \mathcal{A}_g : (X[p^n], \lambda[p^n]) \cong \zeta_k\}.$$

Moreover, $\mathcal{A}_{g, \zeta}$ is affine, smooth and equidimensional.

Proof. In the case $n = \infty$, the existence of a scheme $\mathcal{A}_{g, \zeta}$ satisfying (3.1) is [23, Thm. 3.3]; there, such a scheme is called a “central leaf.” For finite n see, e.g., [27, Thm. 1.2]. It has been understood for some time that $\mathcal{A}_{g, \zeta}$ is pure of codimension $\dim \text{Aut}(\zeta)$ in \mathcal{A}_g [27, 28]. For smoothness, we model our proof on that of [23, Thm. 3.13], which proves the smoothness in the case $n = \infty$. (See also [27, Sec. 4].) Suppose k is algebraically closed and $s, t \in \mathcal{A}_{g, \zeta}(k)$. Choose an isomorphism $(\mathcal{X}, \lambda)_s[p^n] \xrightarrow{\sim} (\mathcal{X}, \lambda)_t[p^n]$; it induces an isomorphism of deformation spaces $\text{Def}((\mathcal{X}, \lambda)_s[p^n]) \xrightarrow{\sim}$

$\text{Def}((\mathcal{X}, \lambda)_t[p^n])$. By [28, Thm. 2.8], $\text{Def}((\mathcal{X}, \lambda)_s[p^\infty]) \rightarrow \text{Def}((\mathcal{X}, \lambda)_s[p^n])$ is formally smooth; the Serre-Tate theorem shows that the formal completion \mathcal{A}_g^s is isomorphic to $\text{Def}((\mathcal{X}, \lambda)_s[p^\infty])$. Of course, the same statements hold at t , *mutatis mutandis*, and taken together this shows that $\mathcal{A}_{g,\xi}^s \subset \mathcal{A}_g^{t/s}$ and $\mathcal{A}_{g,\xi}^t \subset \mathcal{A}_g^{t/s}$ are isomorphic. Since we have given $\mathcal{A}_{g,\xi}$ a structure of reduced subscheme, all of its points are equisingular, thus smooth. \square

The Newton polygon of a BT_∞ is any device which records its geometric isogeny class; the Newton polygon of an abelian variety X is that of $X[p^\infty]$. For an admissible, symmetric Newton polygon ν of height $2g$, let \mathcal{A}_g^ν be the reduced, locally closed subscheme parametrizing those abelian varieties whose Newton polygon is ν .

Oort has shown that \mathcal{A}_g^ν is, up to a finite morphism, a product, as follows.

Theorem 3.2 (Oort). *Let ν be an admissible symmetric Newton polygon of height $2g$. There exist schemes S^ν and T^ν of finite type over \mathbb{F}_p , a geometrically isotrivial p -divisible group $\mathcal{G}^\nu \rightarrow S^\nu \times T^\nu$ with Newton polygon ν , and a finite surjective morphism*

$$(3.2) \quad S^\nu \times T^\nu \xrightarrow{\Phi^\nu} \mathcal{A}_g^\nu$$

such that

- (a) *There is an isogeny $\mathcal{G} \rightarrow \Phi^{\nu*} \mathcal{X}[p^\infty]$.*
- (b) *For each $s \in S^\nu(\mathbb{F})$ there exists some $\xi \in \Xi_{g,\infty}(\mathbb{F})$ such that*

$$\Phi^\nu(s \times T^\nu) = \mathcal{A}_{g,\xi}.$$

- (c) *Each $\mathcal{A}_{g,\xi}$ arises in this way.*

Proof. This is [23, Thm. 5.3]. In fact, what Oort proves is that, over \mathbb{F} , there exists an arrangement as in (3.2) where the domain is a product of *integral* schemes over \mathbb{F} . The present formulation follows by taking the disjoint union of $\text{Gal}(\mathbb{F}_p)$ -conjugates of the schemes constructed in *loc. cit.* \square

Consequently, at least if we set aside the supersingular locus, central leaves have bounded topology. (In general, say that $\xi \in \Xi_{g,n}(\mathbb{F})$ is not supersingular if no component of $\mathcal{A}_{g,\xi}$ is fully contained in the supersingular locus of \mathcal{A}_g .) For a variety V over a field K , let $\sigma_c(V)$ be the sum of its compact Betti numbers;

$$\sigma_c(V) = \sum_i \dim H_c^i(V_{\overline{K}}, \overline{\mathbb{Q}}_\ell).$$

Corollary 3.3. *For each natural number N there exists $C = C(g, N, p)$ such that, for each non-supersingular $\xi \in \Xi_{g,\infty}(\mathbb{F})$, there exists a scheme $V_{\xi,N}$ and a finite and faithfully flat morphism*

$$V_{\xi,N} \xrightarrow{\phi_{\xi,N}} \mathcal{A}_{g,\xi}$$

such that $\phi_{\xi,N}^* \mathcal{X}[p^N]$ is constant and $\sigma_c(V_{\xi,N}) < C$.

Proof. All necessary information is contained in Oort's proof of Theorem 3.2. Subsequent improvements by Chai and Oort [8] allow a somewhat clearer picture of the structure of (3.2), and we avail ourselves of these insights to give a streamlined, if slightly anachronistic, description of Theorem 3.2.

Fix a non-supersingular Newton polygon ν . The uniformization (3.2) is achieved as follows. Let $W = \mathcal{A}_g^\nu$; it is smooth and irreducible. There is an abelian scheme $\mathcal{Y} \rightarrow W$, and an isogeny $\mathcal{Y} \rightarrow \mathcal{X}|_W$ (say of degree p^i), such that \mathcal{Y} admits a polarization of degree p^i and $\mathcal{Y}[p^\infty]$ is completely slope divisible. In particular, the existence of $\mathcal{Y} \rightarrow W$ means there is a map $W \rightarrow \mathcal{A}_{g,p^i}$, whose

image turns out to be a central leaf $C \subset \mathcal{A}_{g,p^i}^\vee$. Fix a central leaf $D \subset W$; there is a finite, surjective morphism $D \rightarrow C$. In (3.2), for T^\vee one takes a finite, faithfully flat $T \rightarrow C$ which trivializes $\mathcal{Y}[p^i]$ (Proposition 3.4) and then sets $T^\vee = T \times_C D$, which is finite and faithfully flat over D . For S^\vee , one takes any isogeny leaf in \mathcal{A}_g^\vee .

Let Y be any fiber of $\mathcal{Y} \rightarrow W$. On one hand, $\mathcal{Y}[p^i]_{T^\vee} = Y[p^i] \times T^\vee$. On the other hand, there is a finite group scheme \mathcal{G}/S^\vee and an inclusion $\mathcal{G} \hookrightarrow Y[p^i] \times S^\vee$ such that $(Y \times S^\vee)/\mathcal{G} \cong \mathcal{X}|_{S^\vee}$ (as polarized abelian schemes). Ultimately, one has

$$\mathcal{G}_{S^\vee \times T^\vee} \hookrightarrow Y[p^i] \times S^\vee \times T^\vee \cong \mathcal{Y}[p^i]_{S^\vee \times T^\vee}.$$

Suppose $s \in S^\vee(\mathbb{F})$; let $\xi = \xi(s)$ be the corresponding isomorphism class of BT_n . The universal abelian scheme over the central leaf $\mathcal{A}_{g,\xi}$ in W is given by

$$\mathcal{X}|_{\mathcal{A}_{g,\xi}} \cong \mathcal{Y}|_{T^\vee} / (\mathcal{G}_s \times T^\vee).$$

In particular, as in Proposition 3.4, let $T^{\vee,N+i} \rightarrow T^\vee$ be a finite and faithfully flat morphism (with $T^{\vee,N+i}$ regular) which trivializes $\mathcal{Y}[p^{N+i}]|_{T^\vee}$. Then for each central stream $\mathcal{A}_{g,\xi} \subset \mathcal{A}_g^\vee$, the finite and surjective morphism $\phi_{\xi,N} : T^{\vee,N+i} \rightarrow \mathcal{A}_{g,\xi}$ trivializes $\mathcal{X}[p^N]|_{\mathcal{A}_{g,\xi}}$. Since the source and target of $\phi_{\xi,N}$ are regular, the morphism is faithfully flat.

The corollary now follows, where for C we may take

$$C = \max \left\{ \sigma_c(T^{\vee,N+i}) : \nu \text{ not supersingular} \right\}.$$

□

In the proof of Corollary 3.3, we needed the fact that a geometrically constant BT_n over a regular base is fppf-trivial, and is in fact trivialized by pullback to a regular scheme:

Proposition 3.4 (Oort). *Let S/K be a regular Noetherian scheme over a perfect field, and let $\mathcal{G} \rightarrow S$ be a geometrically fiberwise constant p -divisible group. For each natural number N , there exists a Noetherian regular scheme T_N/K and a finite and faithfully flat morphism $T_N \rightarrow S$ such that $\mathcal{G} \times_S T_N$ is constant.*

Proof. This is [23, Thm. 1.3]. In fact, Oort shows the existence of a finite, surjective morphism $T_N \rightarrow S$ which trivializes $\mathcal{G}[p^N]$, under the weaker hypothesis that S be Noetherian and integral with Noetherian normalization. Moreover, $T_N \rightarrow S$ factors as $T_N \rightarrow T'_N \rightarrow S_N \rightarrow S$, where $T_N \rightarrow T'_N$ is étale; $T'_N \rightarrow S_N$ is purely inseparable; and the normalization morphism $\tilde{S} \rightarrow S$ factors as $\tilde{S} \rightarrow S_N \rightarrow S$. To prove the formulation in Proposition 3.4, we may assume that S is integral. The normality of S implies that $S_N \cong S$. An analysis of [23, Lemma 1.4] shows that one may realize $T'_N \rightarrow S_N$ as a suitable iterate of the Frobenius morphism $F^{(i)} : T'_N \rightarrow T'^{(p^i)}_N$. Then $T_N \rightarrow S$, as a composition of finite and faithfully flat morphisms, is finite and faithfully flat itself. Finally, suppose $Q \in T'_N$. Since K is perfect, the local rings $\mathcal{O}_{T'_N,Q}$ and $\mathcal{O}_{T'^{(p^i)}_N, F^{(i)}(Q)}$ are (abstractly) isomorphic. Therefore T_N , as an étale cover of a regular scheme, is regular. □

Proposition 3.4 shows that a geometrically fiberwise constant BT_∞ over a regular base is locally constant in the sense of [15]. Let $s \in S$ be a point, with residue field $\kappa(s)$. By [15, Prop. 3], in the setting of Proposition 3.4, for T_n one may take a certain torsor under a finite group scheme over $\kappa(s)$.

Say that $\xi \in \Xi_{g,n}(\mathbb{F})$ admits a generic Newton polygon if there is some Newton stratum $\mathcal{A}_g^{\vee(\xi)}$ such that $\mathcal{A}_g^{\vee(\xi)} \cap \mathcal{A}_{g,\xi}$ is dense in $\mathcal{A}_{g,\xi}$. For example, if $\xi \in \Xi_{g,1}(\mathbb{F})$ is not supersingular, then $\mathcal{A}_{g,\xi}$ is irreducible [11, Thm. 11.5] (see also [16]), and thus ξ admits a generic Newton polygon. At the

opposite extreme, if $n \gg_g 0$, then all pqp p -divisible groups with p^n -torsion ξ are isomorphic, and in particular share the same Newton polygon. More generally, Lau, Nicole and Vasiu give an explicit upper bound [20, Thm. 1.2] for a function (“the isogeny cutoff”) $b(g)$ such that, if $\xi \in \Xi_{g,n}$ and $n \geq b(g)$, then $\mathcal{A}_{g,\xi}$ is contained in a unique Newton stratum. For such a ξ , denote its generic Newton polygon by $\nu(\xi)$.

A stratum $\mathcal{A}_{g,\xi}$ with a generic Newton polygon is irreducible:

Theorem 3.5 (Chai-Oort). *Suppose $1 \leq n \leq \infty$ and that $\xi \in \Xi_{g,n}(\mathbb{F})$ has a generic, non-supersingular Newton polygon. Then $\mathcal{A}_{g,\xi}$ is irreducible.*

Proof. If $n = \infty$ (equivalently, if $n \gg_g 0$), then this is [8, Thm. 4.1]. For arbitrary n , only a very modest adaptation of Chai and Oort’s argument is necessary. It suffices to show that $\mathcal{A}_{g,\xi}^{\nu(\xi)} := \mathcal{A}_{g,\xi} \cap \mathcal{A}_g^{\nu(\xi)}$ is irreducible. Each irreducible component of $\mathcal{A}_{g,\xi}^{\nu(\xi)}$ contains a central leaf, and in particular contains the moduli point of a principally polarized abelian variety isogenous to a hypersymmetric one. Consequently, prime-to- p Hecke correspondences act transitively on the irreducible components of $\mathcal{A}_{g,\xi}^{\nu(\xi)}$. Since $\mathcal{A}_{g,\xi}^{\nu(\xi)}$ is smooth (Lemma 3.1), [7] lets us deduce the irreducibility of $\mathcal{A}_{g,\xi}^{\nu(\xi)}$. \square

4. THE DISTRIBUTION OF FORMS OF GROUP SCHEMES

In all that follows, we will assume that $\xi \in \Xi_{g,n}(\mathbb{F})$ is not supersingular but does admit a generic Newton polygon.

4.1. Generic behavior. Suppose $1 \leq n < m \leq \infty$ and $\xi \in \Xi_{g,n}(\mathbb{F})$. Then there exists at least one lift of ξ to $\Xi_{g,m}(\mathbb{F})$, i.e., there is some $\tilde{\xi} \in \Xi_{g,n}(\mathbb{F})$ such that $\tilde{\xi}[p^n] = \xi$ ([28, Prop. 3.2], which is deduced from the unpolarized version [17, Prop. 1.7]). Then

$$\mathcal{A}_{g,\xi}(\mathbb{F}) = \bigcup_{\tilde{\xi} \in \Xi_{g,m}(\mathbb{F})_{\xi}} \mathcal{A}_{g,\tilde{\xi}}(\mathbb{F}).$$

Moreover, $\bigcup_{\tilde{\xi} \in \Xi_{g,m}(\mathbb{F})_{\xi}} \mathcal{A}_{g,\tilde{\xi}}$ is a disjoint union of countably many locally closed subschemes of \mathcal{A}_g , and is Zariski dense in $\mathcal{A}_{g,\xi}$.

Lemma 4.1. *Fix $1 \leq n < m \leq \infty$.*

- (a) *There exists a subgroup $A_m(\xi) \subseteq A(\xi)$, and an open dense subscheme $\mathcal{A}_{g,m,\xi}^\circ \subseteq \mathcal{A}_{g,\xi}$, such that if $s \in \mathcal{A}_{g,m,\xi}^\circ(\mathbb{F})$, then the image of*

$$(4.1) \quad \text{Aut}_{\mathcal{X}_s[p^m], \lambda_s[p^m]} \longrightarrow \text{Aut}_{\mathcal{X}_s[p^n], \lambda_s[p^n]} \longrightarrow \pi_0(\text{Aut}_{\mathcal{X}_s[p^n], \lambda_s[p^n]}) = A(\tilde{\xi})$$

is $A_m(\xi)$.

- (b) *There is a subset $\Xi_{g,m}^\circ(\mathbb{F})_{\xi} \subseteq \Xi_{g,m}(\mathbb{F})$ such that*

$$\mathcal{A}_{g,m,\xi}^\circ(\mathbb{F}) = \bigcup_{\tilde{\xi} \in \Xi_{g,m}^\circ(\mathbb{F})_{\xi}} \mathcal{A}_{g,\tilde{\xi}}(\mathbb{F}).$$

Proof. Let $\rho : \text{Aut}_{\mathcal{X}[p^m], \lambda[p^m]} \rightarrow \text{Aut}_{\mathcal{X}[p^n], \lambda[p^n]}$ be the obvious homomorphism of group schemes over $\mathcal{A}_{g,\xi}$, and let $K_{m,n}$ be its kernel. Since $\mathcal{A}_{g,\xi}$ is Noetherian, irreducible and reduced, there is a finite decomposition $\mathcal{A}_{g,\xi} = \bigcup_{0 \leq i \leq r} V_i$ of $\mathcal{A}_{g,\xi}$ as a disjoint union of locally closed subschemes such that $K_{m,n}|_{V_i}$ and $\text{Aut}_{\mathcal{X}[p^m], \lambda[p^m]}|_{V_i}$ are flat, and V_0 is open and dense in $\mathcal{A}_{g,\xi}$. Note that, over each V_i , the quotient functor $H := \text{Aut}_{\mathcal{X}[p^m], \lambda[p^m]} / K_{m,n}$ is represented by a group scheme. Possibly after refining the decomposition of $\mathcal{A}_{g,\xi}$, we assume that $H|_{V_i}$ is flat, too.

Since V_0 is Noetherian and $\pi_0(H|_{V_0})$ is constructible, there is an open dense subscheme $U_0 \subseteq V_0$ over which $\pi_0(H)$ has geometrically constant fibers. Let $A_m(\xi)$ be this common group.

For each $1 \leq i \leq r$, let $U_i \subset V_i$ be the (possibly empty) locally closed subscheme where $\pi_0(H) \cong A_m(\xi)$. Then $\mathcal{A}_{g,m,\xi}^\circ := \cup U_i$ is the sought-for open, dense subscheme.

From the definition of $\mathcal{A}_{g,m,\xi}^\circ$ it is clear that, if $s \in \mathcal{A}_{g,\xi}(\mathbb{F})$, then $s \in \mathcal{A}_{g,m,\xi}^\circ(\mathbb{F})$ if and only if (4.1) is surjective. Thus, membership in $\mathcal{A}_{g,m,\xi}^\circ$ depends only on the isomorphism class of $(X[p^m], \lambda[p^m])$, and $\mathcal{A}_{g,m,\xi}^\circ$ is a union of certain locally closed subschemes $\mathcal{A}_{g,\tilde{\xi}}$. \square

Henceforth, let $A_*(\xi) = A_\infty(\xi)$, the group of generically liftable discrete automorphisms of ξ , and let $\mathcal{A}_{g,\xi}^\circ$ be the complement of the supersingular locus in $\mathcal{A}_{g,\infty,\xi}$.

4.2. Distribution of forms. In this section, we fix a $\xi \in \Xi_{g,n}(\mathbb{F})$ which admits a generic, non-supersingular Newton polygon. Recall that $\mathbb{F}_{q(\xi)}$ is the minimal field of definition of a split representative for ξ .

Lemma 4.2.

$$(4.2) \quad \lim_{\mathbb{F}_q/\mathbb{F}_{q(\xi)}} \frac{\#\{(X, \lambda) \in \mathcal{A}_{g,\xi}(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \xi^{\alpha, \mathbb{F}_q} \text{ for some } \alpha \in A_*(\xi)\}}{\#\mathcal{A}_{g,\xi}(\mathbb{F}_q)} = 1.$$

In (4.2) and similar constructions, the limit is over ever-larger extensions of $\mathbb{F}_{q(\xi)}$.

Proof. If the moduli point of (X, λ) is contained in $\mathcal{A}_{g,\xi}^\circ(\mathbb{F}_q)$ (Lemma 4.1), then the isomorphism class of $(X, \lambda)[p^n]$ is represented by a suitable α . \square

Unfortunately, except in certain special cases (such as the *minimal* BT_n 's discussed below), we have little information on the topology of the complement $\mathcal{A}_{g,\xi} \setminus \mathcal{A}_{g,\xi}^\circ$, and thus no explicit control over the rate of convergence in (4.2). Thus, at present, in order to make statements with uniform bounds such as Theorem 4.5, we need to condition on membership in $\mathcal{A}_{g,\xi}^\circ(\mathbb{F}_q)$.

Lemma 4.3. *For each natural number n , there exists $B = B(g, n, p)$ such that the following holds. Let $\tilde{\xi} \in \Xi_{g,\infty}(\mathbb{F})$, and let $\xi = \tilde{\xi}[p^n]$. Then*

$$|A(\tilde{\xi}, \xi)| < B.$$

Proof. Let (G, λ) be a pqp BT_∞ representing $\tilde{\xi}$; it suffices to uniformly bound the size of the (discrete, abstract) group $(\text{End}(G, \lambda)/p^n)^\times$, and thus a uniform bound on $|\text{End}(G)/p^n|^\times$ suffices. Let $m(c/d)$ be the multiplicity of the rational number c/d in the Newton polygon of G . Then $\text{End}(G)$ is an order in $\oplus \text{Mat}_{m(c/d)}(D_{c/d})$, where $D_{c/d}$ is the \mathbb{Q}_p -algebra with invariant c/d , and thus $|\text{End}(G)/p^n|^\times$ admits an upper bound depending only on p, n and the Newton polygon of G . Since, for fixed g , there are only finitely many Newton polygons which arise, the result follows. \square

Lemma 4.4. *Let $\tilde{\xi} \in \Xi_{g,\infty}(\mathbb{F})$ be non-supersingular, and let $\xi = \tilde{\xi}[p^n]$. For each $\alpha \in A(\tilde{\xi}, \xi)$ and each finite extension \mathbb{F}_q of $\mathbb{F}_{q(\xi)}$,*

$$\frac{\#\{(X, \lambda) \in \mathcal{A}_{g,\tilde{\xi}}(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \xi^{\alpha, \mathbb{F}_q}\}}{\#\mathcal{A}_{g,\tilde{\xi}}(\mathbb{F}_q)} = \frac{1}{\#Z_{A(\tilde{\xi}, \xi)}(\alpha)} + O_{g,n,p}(1/\sqrt{q}).$$

In other words, there exists $D = D(g, n, p)$ such that

$$\left| \frac{\#\left\{ (X, \lambda) \in \mathcal{A}_{g, \tilde{\xi}}(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \zeta^{\alpha, \mathbb{F}_q} \right\}}{\#\mathcal{A}_{g, \tilde{\xi}}(\mathbb{F}_q)} - \frac{1}{\#Z_{A(\tilde{\xi}, \xi)}(\alpha)} \right| < \frac{D}{\sqrt{q}}.$$

Proof. Let $s \in \mathcal{A}_{g, \tilde{\xi}}$ be a geometric point. By [8, Thm. 5.6], the naïve monodromy representation

$$\pi_1(\mathcal{A}_{g, \tilde{\xi}}, s) \longrightarrow \text{Aut}(\mathcal{X}_s[p^\infty], \lambda_s[p^\infty]) \cong \text{Aut}(\tilde{\xi})$$

is surjective. Thus, the image of the composite homomorphism

$$\rho_{\tilde{\xi}, \xi} : \pi_1(\mathcal{A}_{g, \tilde{\xi}}, s) \longrightarrow \text{Aut}(\tilde{\xi}) \longrightarrow \text{Aut}(\xi) \longrightarrow \pi_0(\text{Aut}(\xi))$$

is all of $A(\tilde{\xi}, \xi)$. Moreover, the finite and faithfully flat morphism $V_{\tilde{\xi}, n} \rightarrow \mathcal{A}_{g, \tilde{\xi}}$ of Corollary 3.3 trivializes $(\mathcal{X}[p^n], \lambda[p^n])$, and thus the representation $\rho_{\tilde{\xi}, \xi}$; and there is a bound $C = C(g, n, p)$ for the sum $\sigma_c(V_{\tilde{\xi}, n})$ of the ℓ -adic compact Betti numbers of $V_{\tilde{\xi}, n}$. The image $A(\tilde{\xi}, \xi)$ of $\rho_{\tilde{\xi}, \xi}$ has size at most $B = B(g, n, p)$ (Lemma 4.3). The claim now follows from the explicit Chebotarev-type theorem of Katz and Sarnak [18, Thm. 9.7.13]. By [18, Thm. 9.2.6.(4)], for the constant they call $C(\mathcal{A}_{g, \tilde{\xi}}/\text{Spec } \mathbb{F}_{q(\xi)})$, one may take any upper bound for the sum of the compact Betti numbers of an étale Galois cover of $\mathcal{A}_{g, \tilde{\xi}}$ which trivializes $\rho_{\tilde{\xi}, \xi}$; our C provides such a bound. Then Frobenius elements are asymptotically equidistributed in $A(\tilde{\xi}, \xi)$, with error term bounded by $2B \cdot C / \sqrt{q}$. In fact, the Katz-Sarnak estimate only applies when q is at least four times the square of $\sigma_c(\mathcal{A}_{g, \tilde{\xi}})$; but since this may be bounded uniformly in g, n and p , the claimed inequality still holds, possibly at the cost of replacing $2B \cdot C$ with a larger constant D . \square

Theorem 4.5. *Suppose $\xi \in \Xi_{g, n}(\mathbb{F})$ has a generic Newton polygon which is not supersingular. Then for each $\alpha \in A_*(\xi)$ and each finite extension \mathbb{F}_q of $\mathbb{F}_{q(\xi)}$,*

$$\frac{\#\left\{ (X, \lambda) \in \mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \zeta^{\alpha, \mathbb{F}_q} \right\}}{\#\mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q)} = \frac{1}{\#Z_{A_*(\xi)}(\alpha)} + O_{g, n, p}(1/\sqrt{q}).$$

Proof. Decompose $\mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q)$ as a disjoint union

$$\mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q) = \bigcup_{\tilde{\xi} \in \Xi_{g, \infty}^\circ(\mathbb{F})} \mathcal{A}_{g, \tilde{\xi}}(\mathbb{F}_q)$$

and then apply Lemma 4.4:

$$\begin{aligned} \#\left\{ (X, \lambda) \in \mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \zeta^{\alpha, \mathbb{F}_q} \right\} &= \sum_{\tilde{\xi} \in \Xi_{g, \infty}^\circ(\mathbb{F})} \#\left\{ (X, \lambda) \in \mathcal{A}_{g, \tilde{\xi}}(\mathbb{F}_q) : (X, \lambda)[p^n] \cong \zeta^{\alpha, \mathbb{F}_q} \right\} \\ &= \sum_{\tilde{\xi} \in \Xi_{g, \infty}^\circ(\mathbb{F})} \#\mathcal{A}_{g, \tilde{\xi}}(\mathbb{F}_q) \cdot \left(\frac{1}{\#Z_{A_*(\xi)}(\alpha)} + O_{g, n, p}(1/\sqrt{q}) \right) \\ &= \#\mathcal{A}_{g, \xi}^\circ(\mathbb{F}_q) \cdot \left(\frac{1}{\#Z_{A_*(\xi)}(\alpha)} + O_{g, n, p}(1/\sqrt{q}) \right). \end{aligned}$$

\square

5. EXAMPLES AND COMPLEMENTS

We close with some explicit examples of the groups $A(\xi)$ defined in Section 2, and a brief discussion of related work in the literature concerning the distribution of p -power torsion group schemes in abelian varieties.

5.1. Simple minimal p -divisible groups. Fix relatively prime natural numbers c and d ; let $h = c + d$ and let $\mu = d/(c + d)$; and work over any perfect field K containing \mathbb{F}_{p^h} . Then there is a minimal p -divisible group $H_{c,d}/K$ of height h and slope μ . It is characterized by the fact that its endomorphism ring is maximal among all such p -divisible groups; $\text{End}(H_{c,d}) \cong \mathcal{O}_\mu$, the maximal order in the \mathbb{Q}_p -division algebra of Brauer invariant μ [24]. An important fact about minimal p -divisible groups is that they are characterized by their p -torsion; if G is any p -divisible group with $G[p] \cong H_{c,d}[p]$, then $G \cong H_{c,d}$.

On one hand, $\text{Aut}(H_{c,d})$ is the pro-étale group \mathcal{O}_μ^\times . On the other hand, a direct calculation shows that $\pi_0(\text{Aut}(H_{c,d}[p])) \cong \mathbb{F}_{p^h}^\times$.

Rather than prove these statements in complete generality, we work out the special, but representative, case where $(c, d) = (2, 1)$. Then $M := \mathbb{D}_*(H_{2,1})$ admits a $W(K)$ -basis $\{x_1, x_2, x_3\}$ on which the action of Frobenius and Verschiebung are given by

$$\begin{aligned} Fx_1 &= x_2 & Vx_1 &= x_3 \\ Fx_2 &= x_3 & Vx_2 &= px_1 \\ Fx_3 &= px_1 & Vx_3 &= px_2 \end{aligned}$$

Consider an automorphism of α of M as $W(K)$ -module. Then α is an automorphism of M as Dieudonné module if and only if

$$\begin{aligned} A \cdot [F] &= [F] \cdot A^\sigma \\ A^\sigma \cdot [V] &= [V] \cdot A \end{aligned}$$

where $A = (a_{ij})$ is the matrix of α in the chosen basis, and $[F]$ and $[V]$ are similarly given by

$$\begin{aligned} [F] &= \begin{pmatrix} 0 & 0 & p \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \\ [V] &= \begin{pmatrix} 0 & p & 0 \\ 0 & 0 & p \\ 1 & 0 & 0 \end{pmatrix}. \end{aligned}$$

Direct calculation shows that A must be of the form

$$A = \begin{pmatrix} a_{11} & pa_{31}^\sigma & pa_{32}^\sigma \\ a_{32}^{\sigma^2} & a_{11}^\sigma & a_{31}^{\sigma^2} \\ a_{31} & a_{32} & a_{11}^{\sigma^2} \end{pmatrix}$$

where $a_{11}^{\sigma^3} = a_{11}$, i.e., $a_{11} \in W(\mathbb{F}_{p^3})$, and a_{31} and a_{32} similarly lie in $W(\mathbb{F}_{p^3})$.

We now consider the matrix \overline{A} of an automorphism $\overline{\alpha}$ of $M/pM = \mathbb{D}_*(H_{1,2}[p])$. The general form of such an automorphism is

$$\overline{A} = \begin{pmatrix} \overline{a}_{11} & 0 & 0 \\ \overline{a}_{2,1} & \overline{a}_{1,1}^\sigma & 0 \\ \overline{a}_{3,2} & \overline{a}_{2,1}^\sigma & \overline{a}_{1,1}^{\sigma^2} \end{pmatrix}$$

where $\bar{a}_{1,1}^{\sigma^3} = \bar{a}_{1,1}$. Consequently, $\text{Aut}(H_{1,2}[p])$ sits in an exact sequence

$$1 \longrightarrow \text{Aut}(H_{1,2}[p])^0 \longrightarrow \text{Aut}(H_{1,2}[p]) \longrightarrow \mathbb{F}_{p^3}^\times \longrightarrow 1,$$

where the connected component of identity is a two-dimensional unipotent group.

For general (c, d) , one finds that

$$\pi_0(\text{Aut}(H_{c,d}[p])) \cong \mathbb{F}_{p^h}^\times$$

while

$$\pi_0(\text{Aut}(H_{c,d}))/p \cong (\mathcal{O}_\lambda/p)^\times \cong \text{GL}_h(\mathbb{F}_{p^h}).$$

Furthermore, although

$$\text{Aut}(H_{c,d}) \longrightarrow \text{Aut}(H_{c,d}[p])$$

is far from surjective – the source is pro-étale, while the target has positive dimension – the induced map

$$\text{Aut}(H_{c,d}) \longrightarrow \text{Aut}(H_{c,d}[p]) \longrightarrow \pi_0(\text{Aut}(H_{c,d}[p]))$$

is surjective.

5.2. Simple pqp minimal p -divisible groups. Now suppose that $(c, d) \neq (1, 1)$, and let $G_{\mu, 1-\mu} = H_{c,d} \oplus H_{d,c}$. There is a canonical principal $\lambda_{\mu, 1-\mu}$ on $G_{\mu, 1-\mu}$ quasipolarization λ_μ on $H_{c,d} \oplus H_{d,c}$. One has

$$\text{Aut}(G_{\mu, 1-\mu}, \lambda_{\mu, 1-\mu}) \cong \text{Aut}(H_{c,d})$$

and the analogous statement holds for the p -torsion group schemes, as well. Let $\tilde{\xi} \in \Xi_{2h, \infty}(\mathbb{F})$ represent the geometric isomorphism class of $(G_{\mu, 1-\mu}, \lambda_{\mu, 1-\mu})$, and let $\xi \in \Xi_{2h, 1}(\mathbb{F})$ be its p -torsion. Since $G_{\mu, 1-\mu}$ is minimal, we have

$$\mathcal{A}_{2h, \xi} = \mathcal{A}_{2h, \tilde{\xi}} = \mathcal{A}_{2h, \xi}^\circ.$$

Moreover,

$$A(\tilde{\xi}, \xi) = A_*(\xi) \cong \mathbb{F}_{p^h}^\times.$$

5.3. Supersingular elliptic curves. In contrast to Section 5.2, we now consider the case where $c = d = 1$, and restrict our discussion to fields which contain \mathbb{F}_{p^2} . Over an algebraically closed field k of characteristic p , $H_{1,1,k}$ is isomorphic to the p -divisible group of any supersingular elliptic curve over k , and $H_{1,1}[p]_k$ to its p -torsion. Let λ be a principal quasipolarization on $H_{1,1}$. If $a \in \mathbb{F}_{p^2}^\times \subset \text{Aut}(H_{1,1}[p])$, then a scales λ by aa^σ . Consequently,

$$\begin{aligned} \pi_0(\text{Aut}(H_{1,1}[p])) &\cong \mathbb{F}_{p^2}^\times \\ \pi_0(\text{Aut}(H_{1,1}[p], \lambda[p])) &\cong \left\{ a \in \mathbb{F}_{p^2}^\times : aa^\sigma = 1 \right\}. \end{aligned}$$

In particular, we find that $H_{1,1}[p]$ admits $p^2 - 1$ distinct \mathbb{F}_q -forms; as a principally quasipolarized BT_1 , it still admits $p + 1$ distinct \mathbb{F}_q -forms. (Unfortunately, this is not in accordance with [21, Sec. 6]; it seems that, in the cited work, noncanonical identifications of $H_{1,1}[F]$ and its cokernel with α_p led to an incorrect conclusion.)

In fact, following a suggestion of Oort, we can be explicit about these different forms. We work over a perfect field K containing \mathbb{F}_{p^2} . Define a map

$$K^\times \xrightarrow{\tau} K^\times$$

$$a \longmapsto a^\sigma / a^{\sigma^{-1}}.$$

We will exhibit a bijection between forms of $H_{1,1}[p]$ and the quotient $K^\times / \tau(K^\times)$. (Note that $\ker \tau = \mathbb{F}_{p^2}^\times$ and thus, if K is a finite field, $K^\times / \tau(K^\times) \cong \mathbb{F}_{p^2}^\times$.)

Let $\overline{M} = \mathbb{D}_*(H_{1,1}[p])$; since $\dim \operatorname{Hom}(\alpha_p, H_{1,1}[p]) = 1$, \overline{M} is generated as a $K[F, V]$ -module by a single element. (In the coordinates introduced in Section 5.1, the element x_1 suffices.) Moreover, any twist \overline{N} of \overline{M} has this property, too.

For such an \overline{N} and generator x , let $y = Fx$ and define $\lambda_{\overline{N},x} \in K^\times$ by $Vx = \lambda_{\overline{N},x}y$. Any other generator for \overline{N} is of the form $bx + cy$, where $b \in K^\times$ and $c \in K$. For this generator, we find that

$$\begin{aligned} F(bx + cy) &= b^\sigma Fx = b^\sigma y \\ V(bx + cy) &= b^{\sigma^{-1}} \lambda_{\overline{N},x} y \\ &= b^{\sigma^{-1}} (b^\sigma)^{-1} \lambda_{\overline{N},x} F(bx + cy). \end{aligned}$$

Therefore, \overline{N} canonically determines a class $\lambda_{\overline{N}} \in K^\times / \tau(K^\times)$.

Conversely, if $a \in K^\times$, then we can define a form \overline{M}_a of \overline{M} as follows. As a K -vector space, it is generated by x and y ; and the actions of Verschiebung and Frobenius are respectively (since we use covariant Dieudonné theory) given by

$$\begin{aligned} Fx &= y \\ Vx &= ay. \end{aligned}$$

Then $\lambda_{\overline{M}_a,x} = a$, and $\lambda_{\overline{M}_a}$ is the class of a in $K^\times / \tau(K^\times)$.

5.4. The ordinary locus. The techniques developed in the present paper recover known results on the structure of the physical p -torsion of the ordinary locus. Indeed, fix a natural number n and a dimension g ; then the p^n -torsion group scheme of an ordinary abelian variety of dimension g is geometrically isomorphic to $(\mu_{p^n} \oplus \mathbb{Z}/p^n)^{\oplus 2g}$. There is a unique principal quasipolarization on this group scheme, and we let $\zeta_{g,\text{ord}}$ be the corresponding pqp BT $_n$. Then $\operatorname{Aut}(\zeta_{g,\text{ord}}) \cong \operatorname{Aut}((\mathbb{Z}/p^n)^{\oplus 2g}) \cong \operatorname{GL}_g(\mathbb{Z}/p^n)$. Any such automorphism lifts to arbitrarily high level, and $A_*(\zeta_{g,\text{ord}}) \cong \operatorname{GL}_g(\mathbb{Z}/p^n)$. We note, however, that there is a direct proof (Lemma 5.1) of Theorem 4.5 in this case.

5.5. Physical p^n -torsion. For applications (e.g., [6]), it is useful to be able to quantify the distribution of physical p -power torsion. Although Theorem 4.5 can be adapted to such considerations, a direct argument is perhaps more transparent.

Let $X \rightarrow S/\mathbb{F}_{q_0}$ be an abelian scheme of dimension g over a connected base. Suppose that X has constant p -rank $f \geq 1$, in the sense that, for each point $s \in S$, $X_s[p](\overline{\kappa(s)}) \cong (\mathbb{Z}/p)^{\oplus f}$. Then, after a choice of geometric point \bar{s} of S , for each natural number n there is a monodromy representation

$$\pi_1(S, \bar{s}) \xrightarrow{\rho_{X/S,n}} \operatorname{Aut}(X[p^n]^{\text{ét}}) \cong \operatorname{GL}_f(\mathbb{Z}/p^n).$$

If $\mathbb{F}_q/\mathbb{F}_{q_0}$ is any extension and if $t \in S(\mathbb{F}_q)$, then there is an element $\sigma_{t/\mathbb{F}_q} = \rho_{X/S,n}(\text{Fr}_t) \in \text{GL}_f(\mathbb{Z}/p^n)$, well-defined up to conjugacy. The q -power map $a \mapsto a^q$ acts on $X_t[p^n](\overline{\mathbb{F}}_q)$ as σ_{t/\mathbb{F}_q} , and this action determines the isomorphism class of $X_t[p^n]^{\text{ét}}$.

Lemma 5.1. *With notation as above, suppose that $\rho_{X/S,n}$ is surjective. For each $\alpha \in \text{GL}_f(\mathbb{Z}/p^n)$ and each $\mathbb{F}_q/\mathbb{F}_{q_0}$,*

$$(5.1) \quad \frac{\#\left\{t \in S(\mathbb{F}_q) : \sigma_{t/\mathbb{F}_q} \sim \alpha\right\}}{\#S(\mathbb{F}_q)} = \frac{1}{\#Z_{\text{GL}_f(\mathbb{Z}/p^n)}(\alpha)} + O_{X/S}(1/\sqrt{q}).$$

Proof. This is presumably well-known. However, since there seems to be some uncertainty in the literature (e.g., the discussion surrounding Principle 3 in [6]), we provide a sketch. The étale sheaf

$$\mathcal{I} := \text{Isom}(X[p^n]^{\text{ét}}, (\mathbb{Z}/p^n)^{\oplus f})$$

on S is represented by a scheme. The scheme \mathcal{I} is visibly an étale cover of S , with geometric fiber $\text{GL}_f(\mathbb{Z}/p^n)$; the hypothesis that $\rho_{X/S,n}$ is surjective is equivalent to the irreducibility of \mathcal{I} . The claimed equidistribution now follows from the Chebotarev theorem for schemes [25, Thm. 7]. \square

In many applications, there exists a (partial) compactification $X \rightarrow T$ of the family, and the étale cover $\mathcal{I} \rightarrow S$ extends to a wildly ramified cover of T . Although some of the explicit estimates of [18] need not apply – see, for example, the discussion in [19, Rem. 4.8] – there is always a suitable Chebotarev-type theorem, with error term $O(1/\sqrt{q})$. In the context of Lemma 5.1, the apparent counterexamples in [19] stem from a naïve attempt to estimate the Betti numbers of \mathcal{I} in terms of the Betti numbers of S and the degree of $\text{GL}_f(\mathbb{Z}/p^n)$. The sum of the compact Betti numbers of \mathcal{I} , whatever it might be, is the input to an explicit error term in (5.1).

Let $\mathcal{A}_g^f \subset \mathcal{A}_g$ be the moduli space of principally polarized abelian varieties of p -rank f . The surjectivity of each $\rho_{X/\mathcal{A}_g^f,n}$ is well-known [10, 13]. For $1 \leq f \leq g-1$, the surjectivity of $\rho_{X/\mathcal{A}_g^f,n}$ follows from the surjectivity of $\rho_{X/\mathcal{A}_f^f,n}$, the existence of the natural inclusion $\mathcal{A}_f^f \times \mathcal{A}_{g-f}^0 \hookrightarrow \mathcal{A}_g^f$, and the irreducibility of \mathcal{A}_g^f .

5.6. The torsion of Jacobians. If C/\mathbb{F}_q is a smooth, projective curve, then the class group of its function field is isomorphic to $\text{Jac}(C)(\mathbb{F}_q)$. Consequently, results about the distribution of torsion on (hyperelliptic) *Jacobians* have implications for Cohen-Lenstra type results for function fields. For prime-to- p torsion see, e.g., [1] for fixed g , large q results; and [12] for the much more difficult case of fixed q and large g .

For physical p -power torsion, and thus for elements of order p in the class groups of (hyperelliptic) function fields, all the results surveyed in Sections 5.4 and 5.5 have immediate parallels for \mathcal{M}_g^f , the moduli space of hyperelliptic curves of genus g and p -rank f , and for the hyperelliptic sublocus $\mathcal{H}_g^f \subset \mathcal{M}_g^f$. Indeed, for each irreducible component S of \mathcal{M}_g^f or \mathcal{H}_g^f , one knows that the monodromy representation $\pi_1(S) \rightarrow \text{GL}_f(\mathbb{Z}/p^n)$ is surjective [2, 3]. Thus, Lemma 5.1 computes the distribution of physical p -torsion for (hyperelliptic) Jacobians.

However, it is not clear how to apply the methods of the present paper to strata $\mathcal{M}_{g,\xi}$ or $\mathcal{H}_{g,\xi}$. Indeed, it is not even known which Newton strata in \mathcal{M}_g are nonempty – see [4] for a survey of what little is known – let alone finer strata $\mathcal{M}_{g,\xi}$, to say nothing of their p -adic monodromy. Thus, a version of Theorem 4.5 for hyperelliptic Jacobians remains a distant goal.

REFERENCES

- [1] J. D. Achter. Results of Cohen-Lenstra type for quadratic function fields. In *Computational arithmetic geometry*, volume 463 of *Contemp. Math.*, pages 1–7. Amer. Math. Soc., Providence, RI, 2008.
- [2] J. D. Achter and R. Pries. Monodromy of the p -rank strata of the moduli space of curves. *Int. Math. Res. Not. IMRN*, (15):Art. ID rnn053, 25, 2008.
- [3] J. D. Achter and R. Pries. The p -rank strata of the moduli space of hyperelliptic curves. *Adv. Math.*, 227(5):1846–1872, 2011.
- [4] J. D. Achter and R. Pries. Generic Newton polygons for curves of given p -rank. In *Algebraic Curves and Finite Fields: Cryptography and Other Applications*, volume 16 of *Radon Series on Computational and Applied Mathematics*, pages 1–22. de Gruyter, Boston and Berlin, 2014.
- [5] B. Cais, J. S. Ellenberg, and D. Zureick-Brown. Random Dieudonné modules, random p -divisible groups, and random curves over finite fields. *J. Inst. Math. Jussieu*, 12(3):651–676, 2013.
- [6] W. Castryck, A. Folsom, H. Hubrechts, and A. V. Sutherland. The probability that the number of points on the Jacobian of a genus 2 curve is prime. *Proc. Lond. Math. Soc.* (3), 104(6):1235–1270, 2012.
- [7] C.-L. Chai. Monodromy of Hecke-invariant subvarieties. *Pure Appl. Math. Q.*, 1(2):291–303, 2005.
- [8] C.-L. Chai and F. Oort. Monodromy and irreducibility of leaves. *Ann. of Math.* (2), 173(3):1359–1396, 2011.
- [9] H. Cohen and H. W. Lenstra, Jr. Heuristics on class groups of number fields. In *Number theory, Noordwijkerhout 1983 (Noordwijkerhout, 1983)*, volume 1068 of *Lecture Notes in Math.*, pages 33–62. Springer, Berlin, 1984.
- [10] T. Ekedahl. The action of monodromy on torsion points of Jacobians. In *Arithmetic algebraic geometry (Texel, 1989)*, pages 41–49. Birkhäuser Boston, Boston, MA, 1991.
- [11] T. Ekedahl and G. van der Geer. Cycle classes of the E-O stratification on the moduli of abelian varieties. In *Algebra, arithmetic, and geometry: in honor of Yu. I. Manin. Vol. I*, volume 269 of *Progr. Math.*, pages 567–636. Birkhäuser Boston Inc., Boston, MA, 2009.
- [12] J. Ellenberg, A. Venkatesh, and C. Westerland. Homological stability for Hurwitz spaces and the Cohen-Lenstra conjecture over function fields, December 2009, arXiv:0912.0325. arXiv:0912.0325.
- [13] G. Faltings and C.-L. Chai. *Degeneration of abelian varieties*. Springer-Verlag, Berlin, 1990. With an appendix by David Mumford.
- [14] O. Gabber and A. Vasiu. Dimensions of group schemes of automorphisms of truncated Barsotti-Tate groups. *Int. Math. Res. Not. IMRN*, (18):4285–4333, 2013.
- [15] M. A. Garuti. Barsotti-Tate groups and p -adic representations of the fundamental group scheme. *Math. Ann.*, 341(3):603–622, 2008.
- [16] S. Harashita. Generic Newton polygons of Ekedahl-Oort strata: Oort’s conjecture. *Ann. Inst. Fourier (Grenoble)*, 60(5):1787–1830, 2010.
- [17] L. Illusie. Déformations de groupes de Barsotti-Tate (d’après A. Grothendieck). *Astérisque*, (127):151–198, 1985. Seminar on arithmetic bundles: the Mordell conjecture (Paris, 1983/84).
- [18] N. M. Katz and P. Sarnak. *Random matrices, Frobenius eigenvalues, and monodromy*. American Mathematical Society, Providence, RI, 1999.
- [19] E. Kowalski. The large sieve, monodromy and zeta functions of curves. *J. Reine Angew. Math.*, 601:29–69, 2006.
- [20] E. Lau, M.-H. Nicole, and A. Vasiu. Stratifications of Newton polygon strata and Traverso’s conjectures for p -divisible groups. *Ann. of Math.* (2), 178(3):789–834, 2013.
- [21] C. Liedtke. The p -torsion subgroup scheme of an elliptic curve. *J. Number Theory*, 131(11):2064–2077, 2011.
- [22] B. Moonen. Group schemes with additional structures and Weyl group cosets. In *Moduli of abelian varieties (Texel Island, 1999)*, pages 255–298. Birkhäuser, Basel, 2001.
- [23] F. Oort. Foliations in moduli spaces of abelian varieties. *J. Amer. Math. Soc.*, 17(2):267–296 (electronic), 2004.
- [24] F. Oort. Minimal p -divisible groups. *Ann. of Math.* (2), 161(2):1021–1036, 2005.
- [25] J.-P. Serre. Zeta and L functions. In *Arithmetical Algebraic Geometry (Proc. Conf. Purdue Univ., 1963)*, pages 82–92. Harper & Row, New York, 1965.
- [26] J.-P. Serre. *Galois cohomology*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, english edition, 2002. Translated from the French by Patrick Ion and revised by the author.
- [27] A. Vasiu. Level m stratifications of versal deformations of p -divisible groups. *J. Algebraic Geom.*, 17(4):599–641, 2008.
- [28] T. Wedhorn. The dimension of Oort strata of Shimura varieties of PEL-type. In *Moduli of abelian varieties (Texel Island, 1999)*, pages 441–471. Birkhäuser, Basel, 2001.

E-mail address: j.achter@colostate.edu

DEPARTMENT OF MATHEMATICS, COLORADO STATE UNIVERSITY, FORT COLLINS, CO 80523

URL: <http://www.math.colostate.edu/~achter>